Mobile Robot 3D Perception and Mapping with Multi-Resolution Occupancy Lists

Julian Ryde and Huosheng Hu

Department of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, England {jryde, hhu}@essex.ac.uk

Abstract—Many real-world applications require mobile robots to be able to implement 3D perception and mapping. This paper proposes a novel mechanism for augmenting a traditional 2D laser range finder to produce 3D scans. The range data is stored in occupancy lists which are aligned to produce 3D maps by a multi-resolution particle filter. Experimental results are presented to show the feasibility and good performance of the proposed approach.

Index Terms—Particle filter, Localization, Mapping, 3D laser scanner

I. INTRODUCTION

Extracting 3D information about the world surrounding a mobile robot has proved difficult. The two main approaches, vision and laser range finding, have been beset by problems. Vision is often computationally intensive and suffers from sensitivity to changes in illumination. Many of the difficulties stem from having to solve the correspondence problem which can be alleviated by structured light approaches, however the data spatial density does not come close to that provided by laser scanners. Non-visual localization and mapping has taken place in 2D, mainly due to limitations of the sensors in the case of laser range finders, or processor speed and algorithms for that of stereoscopic vision.

Recently in a drive to test the benefits of 3D sensors researchers have mounted 2D laser scanners on nodding or rotating mechanisms to obtain 3D scans [6], [8], [9]. Alternatively, two laser scanners mounted with their scan planes orthogonal [4] are also popular. It must be noted that the orthogonal mounting technique can produce 3D maps however it does not give the robot 3D sensory perception which is necessary for reliable obstacle avoidance. Significant attention is being focused on 3D laser mapping [2], [3], [7].

A completely different approach is undertaken by [13]. This scanner has a modulated infra-red illumination and a CMOS/BCCD image sensor taking four samples of the reflected infra-red illumination intensity to determine the phase difference of the reflected data for each image pixel. From this phase difference and the frequency of illumination modulation the distance to image pixels can be ascertained unambiguously up to 7.5m. This sensor delivers 19480 range measurements at 30 FPS, with no moving parts, whilst comparing favourably in accuracy to the SICK LMS 200. Its main disadvantages are price, low availability and a relatively narrow field of view, although the latter can be altered by the lens arrangement.

There are numerous commercial 3D laser range finders however these are often prohibitively expensive especially when required for multiple robots. They tend to be designed for surveying and so often have a narrow field of view and sacrifice scan rate for high point density.

This multitude of approaches illustrates the rapidly increasing interest in full 3D robotic sensory perception. The detection of negative and over-hanging obstacles greatly enhances avoidance behaviour. Once 3D maps of environments have been built they can be customized for and distributed to different robots. For instance various 2D occupancy grids may be built for robots of different sizes or with 2D sensors at different heights. Severely cluttered non-manifold environments such as search and rescue situations may be reliably mapped. Maps based upon the ceilings of rooms [12] will remain accurate for longer and a clear view of the ceiling is usually readily accessible to a robot even in crowded environments [1]. An alternative approach taking advantage of full 3D perception [14] uses virtual 2D scans produced by collapsing the 3D data vertically into a plane and then taking the furthest point for each 2D scan angle. This produces good 2D representations in cluttered environments. The disadvantages of 3D sensing technologies are slower scan acquisition time and the geometric increase in data to be processed.

Section II contains a description of the hardware improvements that produce 3D scans from a laser scanner. The mechanism for scan matching to produce maps and global localisation is explained in Section III. A theoretical analysis of the significance levels of overlap counts at varying resolutions is undertaken in Section IV. Section V displays the experimental results for assessing the accuracy of the 3D laser scanner, the distribution of the overlap counts and map building with the multi-resolution particle filter. The paper ends with a summary and discussion of future work in Section VI.

II. ENHANCED 3D LASER SCANNER

We have developed a novel 3D laser scanner based on a SICK LMS 200, [5], [10], [11] which consists of the LMS 200 facing upwards into a rotating mirror driven by a stepper motor. This approach simplifies the hardware and software implementation whilst producing a 3D laser scanner suitable for real-time operation on medium sized (≈ 1 m) mobile robots. The high update rate (75Hz) of the LMS 200



Fig. 1. Pioneer 2 robot with rotating mirror mechanism enabling 3D scans.

means the 3D scanner delivers scans at 1Hz with a horizontal resolution of 1° and vertical resolution of 5° .

The processing of data produced by the scanner is similar to that described in [10], however it is substantially simpler to implement due to the accuracy of the open loop control provided by the stepper motor driven by a frequency divided quartz crystal oscillator. Stepper motors can be operated over a range of speeds and are especially suited to low speed operation. Directly coupling the motor to the mirror eliminates backlash by removing the gear train that introduces backlash. As long as the stepper motor's maximum load torque is not exceeded its rotation speed is solely determined by the frequency of pulses it receives from the driver circuit. Variations in load torque do not affect its rotation speed and so closed loop control is no longer necessary and substantial hardware and software implementation complexity is removed.

The stepper motor in this implementation was a 1.8° resolution that was half-stepped to give 400 steps per revolution. Although the stepper motor driven by a quartz crystal oscillator has an exceedingly consistent angular velocity the mirror still needs to be approximately balanced. Significant off-centre mass distribution will cause flexing to occur between the motor connection and the mirror. This flexing may be reduced by using stiffer materials and balancing the mirror to reduce the variation in load torque. The system however is much less sensitive to variations in load torque than a DC motor driven system. A 12MHz oscillator is stepped down to produce a highly stable and accurate 15ms clock signal.

The field of view is improved by placing the laser scanner facing upwards. In this way the blind spot when the mirror is edge on is rotated to point upwards rather than in front of the robot. The observable volume is given by the following constraints.



Fig. 2. Top and side views of the enhanced 3D laser scanner.

$$0.5m < r < 8m, 25^{\circ} < \theta < 155^{\circ}, 75^{\circ} < \phi < 105^{\circ}.$$

Data from a single scan of a small room is plotted in Fig. 8. The room is a cuboid and with flat ceiling, walls and floor.

The SICK LMS 200 is capable of operating at higher data rates. The internal mirror rotates at 75Hz and this is the natural data frequency if the RS232 serial communications bottle-neck is removed. The SICK LMS supports a high-speed serial connection RS422 which may be connected to a standard USB port with a USB-RS422 converter. This allows 180° of range data at 1° resolution to be delivered to the host computer at 75Hz. The horizontal resolution is fixed at 1° however the vertical resolution may be adjusted by varying the speed of the external mirror rotation. Matching the vertical resolution to the horizontal resolution gives scan times of 360/75 = 4.8 seconds. The resulting scans are high detail which can be used for map building, however for obstacle avoidance a faster mirror speed is more suitable and scan frequencies of 1Hz are feasible.

Fig. 2 indicates the positioning of the blocking arms. By blocking the range scan every half mirror rotation these arms serve two main purposes. The first is to establish the angular velocity of the mirror, although it is very consistent measuring the angular velocity makes the system more flexible and aids fault detection. The second is to establish the rotation of the scan about the *y*-axis so that the resulting 3D scan is correctly orientated. This does not have to be too accurate as this is affected by the inclination of the surface that the robot is traversing. In this manner the system does have an element of closed-loop feedback once per cycle but most of the fine control is an open-loop system thus significantly reducing demands on the host hardware.

If the reflecting surface used is a standard single-sided mirror then only one blocking arm should attached otherwise there will be a 180° ambiguity in ϕ . With a double-sided mirror and two blocking arms two 3D scans are returned for every mirror revolution.

III. MULTI-RESOLUTION PARTICLE FILTER

Global localisation and scan matching are achieved by an approach that is similar to a multi-resolution particle filter. In this work a particle is considered as a pose with an associated probability weight w. In the implementation



Fig. 3. Flowchart of the particle filter global localisation and alignment process.

these weights are not normalised and are integers and are considered proportional to the likelihood of the corresponding pose.

Data: OccupiedList map, OccupiedList scan, Integer f**Result**: Set of weighted poses

generate uniform pose distribution set particleSet;

for f is 2^i where i is 6 to 0 decrement 1 do

down-sample map and local scan by f;

calculate overlaps for all poses in particleSet;

remove from particleSet all poses with certainty <

80% of the mode;

re-sample poses around remaining poses in

particleSet uniformly with width $\epsilon/2$ spatially and $\Delta\theta/2$:

end

Algorithm 1: Multi-resolution particle filter overview algorithm

In Algorithm 1 the re-sampling step is akin to sampling importance re-sampling with a local uniform distribution. Typically in particle filters the initial distribution is randomly generated from a uniform distribution. In this implementation the initial pose distribution is exhaustively generated for all possible initial poses. This is made computationally possible by the down-sampling step and the overlap calculation algorithm which ensures the smoothing of the probability distribution function at low resolutions.

The implementation of the particle set must satisfy two



Fig. 4. Distribution of best orientation poses with $w > 0.8w_{\text{max}}$ as a function of 2D position with height indicating probability of pose.

criteria. It must only store unique poses at the current spatial and angular resolutions and it must be able to iterate through the particle poses in angular order. The latter is not a requirement for success of the process but is for optimization reasons. When calculating the overlaps one of the slowest operations is the rotation step. By ordering the poses particles by rotation the rotation on the local scan may be done only once and then translated to poses at different positions with identical orientation. This is implemented by using a treeset structure which maintains the uniqueness and correct order of elements and guarantees insertion times of $O(\log n)$.

A typical particle distribution is visualised in Fig. 4 where the local scan (depicted in 3D) is globally aligned with the map (projected onto the floor). Both the local scan and the map are full 3D representations however they are shown as 2D projections in Fig. 4 for clarity. The map was built by combining 20 scans with the multi-resolution particle filter and then another scan was globally localised within this map. The initial particle distribution is uniform with the spatial separation of 1.24m. The particles are re-sampled around the highest probability regions progressively sampling the pose space at better resolutions. The final alignment is at a resolution of 0.02m.

The solution is indicated by the high cluster of poses shown in Fig. 4. The algorithm is similar for global localisation as well as tracking with the only difference being the first step of generating the initial particle distribution. For global localisation the particle distribution is over the entire possible pose space. The two main ways of generating the initial particle distribution are uniform and random. In other work with particle filters the random initialisation is usually preferred. For this algorithm it is important to use a uniform initial particle distribution with the spatial and angular resolution of the poses carefully chosen. In this way it is possible to be sure of complete coverage of the pose space. Normally this would be exceedingly costly however down-sampling makes this possible. Other advantages of uniform initial particle distribution are that it is deterministic and quick to generate.

The spatial Δx and angular $\Delta \theta$ resolution need to be



Fig. 5. The quantisation process for converting pointclouds to the occupied voxel lattice. The original pointcloud data is represented by crosses. Two different resolution grids are also shown.

matched and are dependent on the resolution ϵ of the occupancy lists and the scan range, R. The spatial separation of the particles is dictated by the resolution of occupancy list being processed. $\Delta x = \epsilon$ The angular separation is

$$\Delta \theta = \epsilon / R,\tag{1}$$

where R is the maximum range of points from the origin of the local map which need to be reliably considered. This ensures that successive rotations produce new local occupancy lists at the current resolution. Because the particles are stored in sets then the random initial particle distribution is still guaranteed to produce unique particles.

An important aspect of this process is the down-sampling of the occupancy lists. The 3D laser scanner returns 3D scans as point clouds and these are converted into an occupancy list of resolution $\epsilon = 0.02m$. The resolution of 0.02m is chosen because this is the limit of the accuracy of the 3D laser scanner. The down-sampling process is described in Algorithm 2 and Fig. 5.

Data: occupanyList and factor
Result: lower resolution occupiedList
if *factor exists in downSample cache* then return corresponding downSampled occupiedList;
end
make new occupiedList with resolution of factor*this resolution;
for *all voxels in occupiedList* do add new voxel(Math.round(v.x/f), Math.round(v.y/f),

end

Math.round(v.z/f));

store factor with downSampled occupiedList; Algorithm 2: Occupany list down-sample algorithm.

Once the occupancy lists have been down-sampled and the particles selected the probabilities or weights for each particle need to be estimated. This process is finely balanced between giving an accurate weighting and computational speed as it is carried out for each particle. The process is described in Algorithm 3 and my be succinctly expressed as

$$o = |L \cap M|. \tag{2}$$

The overlap count is o, the number of elements or cardinality of set A is |A|, the map is M and the local transformed scan is L. Thus o is not only a count of the overlapping voxels but also an indicator of the probability of the associated pose particle. A statistical analysis of the significance levels for oin various situations is presented in Section IV.

Data: particleSet poseParticles, occupiedList map, occupiedList local Result: particleSet poseParticles with updated weights

for each particle in poseParticles do o = 0:

if *particle angle is different from previous* then calculate new rotated local *L*;

end translate L by (p_x, p_y) ; for each element in L do if element is present in M then increment o; end end

set particle weight to o; if $o > w_{max}$ then $o_{max} = o$; end

end

Algorithm 3: Process for updating the weights of the particle set.

It is assumed that |L| < |M| and so iteration is over the elements of L. The particles are stored in a set sorted primarily by their angle. Rotation operations on the local occupancy list are slower than translation operations. Rotations in 2D Cartesian coordinates require four multiplications and two additions as opposed to translations which require two additions. Sorting the list by rotation groups identical rotations together hence ensuring rotations need to be performed once only.

Once the particle weights have been updated, particles with a weight less than 80% of the maximum weight are removed. New particles are then generated around the remaining highweight particles as described in Algorithm 4.

This new particle set is then submitted for the next round of tests at better resolution. In this manner an order of 10,000 particles per second can be assessed on a standard 2GHz computer. The precise rate depends on the resolution and structure of the map and local occupancy lists.

IV. SCAN MATCHING PROBABILITY

Whilst aligning a scan with the map it is important to determine the significance of a particular overlap value in order to assess the reliability of that localisation. Determination of the confidence in the alignment pose is required for reliable map building; only those scans with high confidence value Data: particleSet initialParticles **Result**: particleSet with re-sampled particles integers a, b, c; create new particleSet P; for each particle in initialParticles do for $-1 \le a \le 1$ do for $-1 \le b \le 1$ do for $-1 \le c \le 1$ do create a particle with pose $(x + a\epsilon/2, y + b\epsilon/2, \theta + c\Delta\theta);$ add particle to P; end end end end return P;

Algorithm 4: Re-sampling at improved spatial and angular resolution around high weight particles. This algorithm will produce repeated particles for adjacent initial particles however particles are placed in a set to ensure uniqueness.

should be added to the map. The shape of the probability distribution is useful in establishing this significance. For instance if the distribution is multi-modal and the peaks are similar in height then there is doubt as to which modal peak is correct. Likewise the sharpness of the peaks indicates the accuracy of the corresponding modal poses. These measures are all relative; they say nothing about the absolute value of the overlap. This is addressed by the following analysis. In order to determine significance of a given overlap the probability of getting an overlap at least that high by pure chance alone has to be considered.

Let the volume covered by the map be M_v , the overlap o, the number of occupied voxels in the map and scan is N and n respectively. The map and scan voxel sizes are both ϵ . The total number of possible voxels in the map volume is V/ϵ^3 . Thus the proportion of the map that is occupied is

$$\frac{M\epsilon^3}{M_v} \tag{3}$$

Assuming a random distribution of occupied voxels throughout the map and that voxel occupancy is independent, Equation 3 expresses the probability of a single voxel picked at random being occupied in the map. Given a scan of N voxels what is the expected value of the overlap assuming that the scan voxel occupancy probabilities are independent? This can be re-expressed in standard probability terminology as picking n voxels at random from the map without replacement. This means that the overlap counts may be modelled as a hypergeometric distribution. The mean of which is given by

$$\bar{x} = \frac{Mo}{N}.$$
(4)

The variance is

$$\sigma^2 = \frac{n(o/N)(1 - o/N)(N - n)}{N - 1}.$$
(5)



Fig. 6. Probability distribution of overlap counts for various resolutions of the map and local scan.

It is important to note that for simplicity the hypergeometric distribution tends towards a normal distribution as N >> n. The significance of the maximum overlap value may be determined by calculating its probability of occurrence.

The actual distributions of overlaps for a number of resolutions are shown in Fig. 6 which shows that the mean overlap stays relatively constant regardless of resolution however the standard deviation grows as the resolution becomes finer. What is also clear is that the overlap count may be adequately represented by a normal distribution down to resolution of 0.64m. The region of interest is the far right of this graph where the significance of the maximum overlap may be established. The maximum count for the 0.16 resolution scan alignment shown in Fig. 4 is 379 which is very significant and so the result may be regarded as reliable. It must be recalled that one of the original assumptions was the independence of voxel occupancy. This assumption becomes more valid as ϵ , the voxel size, increases.

V. SCANNER ERROR

The random error for the 3D scanning mechanism was found by scanning a room with flat ceiling and walls, Fig. 8. A subset of the data points associated with the ceiling which was 3.2m above the origin of the 3D laser scanner was extracted from the scan. This extract indicates that the error in z is approximately 0.02m. Repeating this analysis for the smooth walls of the room indicate a similar error in x and y. The error is relatively independent of the distance over the range of distances 0-8m implying that a simplified sensor model is appropriate. In this simplified sensor model the errors are 0.02m regardless of distance. This also places a lower bound on the map resolution.

To establish the size of the systematic error numerous dimensions of the experimental environment are compared to the generated map and the mean error is 0.07m on dimensions of 6m. The systematic error results would seem to indicate that there is no point using map resolutions better than 0.07m and experimentally we find the best localization and mapping performance with resolutions of 0.08m, however it is worth remembering that local regions of the scans are accurate to 0.02m.



Fig. 7. Distribution of the distances to the corresponding closest points in two sequential scans.



Fig. 8. Side view of a single scan for a room of dimensions 3 by 6 by 2.75. The grid cells are 1m.

Having established an estimate for the systematic error of a scan the random error is now inspected. This is done by analysing two consecutive scans in a static environment. Ideally the scans should be identical because nothing in the environment has changed between the scan times. The differences between the two scans are due to random error as any systematic distortions will be equally present in both scans. For each point in the first scan the distance to the closest point in the second scan is recorded. The distribution of these corresponding closest point distances is graphed in Fig. 7.

Although the mean closest point separation is 0.011m the median is 0.006m and the mode is 0.004m. The random error is remarkably small with very few separations venturing above 0.02m. Thus it may be concluded that the majority of the error comes from systematic errors or distortions within the laser scan. In this case flexing between the stepper motor axle and the mirror. From the point of view of mobile robotics the data accuracy is significantly better than that acquired by stereoscopic vision and comparable to that acquired by rotating the laser scanner. For medium sized robots in standard human environments this 0.02m accuracy is adequate and paves the way for equally accurate mapping and localisation.

VI. CONCLUSION AND FUTURE WORK

This work describes the enhancements made to a standard SICK LMS 200 laser scanner to produce a full 3D sensor. This sensor has medium power consumption, minimal hardware requirements and may be easily deployed on standard Pioneer 2 and other robotics platforms with little or no modification. This was achieved by mounting a rotating mirror above an upward facing 2D scanner to divert the 2D scans to gather range data from a volume rather than a single scan plane. Knowledge of the mirror's angular velocity and speed allows this range data to be converted into 3D point clouds that contain information regarding the obstacles surrounding the robot. The addition of the external rotating mirror has a negligible impact on the accuracy of the 2D laser scanner. The 3D scans have an error of 0.02m, consist of up to 5000 data points with some blind spots and are produced at 1Hz.

The scans are fused with a multi-resolution particle filter to produce 3D maps of the environment which allow localisation to $\pm 0.02m$ and sub-degree error.

Further work will include scaling up the experiments to larger areas and the inclusion of pose corrections when loop closing has been detected. The mapping and localisation performance in dynamic environments will be assessed.

References

- W. Burgard, A. B. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1–2):3–55, 1999.
- [2] D. Hähnel, D. Schulz, and W. Burgard. Map building with mobile robots in populated environments. In Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2002.
- [3] A. Howard, A. Parker, E. Lynne, and G. S. Sukhatme. The SDR experience: Experiments with a large-scale heterogenous mobile robot team (extended abstract). *Int. Symposium on Experimental Robotics*, June 2004.
- [4] A. Howard, D. F. Wolf, and G. S. Sukhatme. Towards 3D mapping in large urban environments. In *IEEE/RSJ Int. Conf. on Intelligent Robots* and Systems, pages 419–424, Sendai, Japan, Sept. 2004.
- [5] J. Ryde and H. Hu. Laser based simultaneous mutual localisation for multiple mobile robots. In *Proc. of Int. Conf. Mechatronics and Automation*, pages 404–409, Niagara Falls, Canada, July 2005.
- [6] K. Lingemann, H. Surmann, A. Nüchter, and J. Hertzberg. Indoor and outdoor localization for fast mobile robots. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 2185–2190, Sendai, Japan, Sept. 2004.
- [7] M. Montemerlo and S. Thrun. Large-scale robotic 3D mapping of urban structures. In Proc. of the Int. Symp. on Experimetnal Roboics (ISER), 2004.
- [8] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann. Heuristicbased laser scan matching for outdoor 6D SLAM. In Advances in artificial intelligence. 28th annual German Conf. on AI, Sept. 2005.
- [9] P. Pfaff, R. Triebel, and W. Burgard. An efficient extension to elevation maps for outdoor terrain mapping and loop closing. *The Int. Journal* of Robotics Research, 26(2):217–230, 2007.
- [10] J. Ryde and H. Hu. Cooperative mutual 3D laser mapping and localization. In *IEEE International Conference on Robotics and Biomimetics*, Kunning, China, 12 2006.
- [11] J. Ryde and H. Hu. Mutual localization and 3D mapping by cooperative mobile robots. In *Intelligent Autonomous Systems*, The University of Tokyo, Tokyo, Japan, 7 Mar. 2006.
- [12] W. Jeong and K. Lee. CV-SLAM: A new Ceiling Vision-based SLAM technique. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2005.
- [13] J. Weingarten, G. Gruener, and R. Siegwart. A state-of-the-art 3D sensor for robot navigation. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Oct. 2004.
- [14] O. Wulf, K. O. Arras, H. I. Christensen, and B. Wagner. 2D mapping of cluttered indoor environments by means of 3D perception. In *IEEE/RAS Int. Conf. on Robotics and Automation (ICRA)*, New Orleans, USA, 26 Apr. 2004.